

Perspectives on the Memory Wall

“It’s the Memory, Stupid!”

— Richard Sites

First, My Philosophy...

Use existing resources more wisely

Add **minimal** hardware support, **isolate** complexity

Modify software (OS/compiler/libs/apps) to exploit that hardware



The Game Plan

What's the problem?

- Numbers
- Pictures
- Details, details

What are we going to do about it?

- Good news
- Bad news
- Silver Bullet?

(Selective, Subjective) Chronology

1995: Mark G. Karpman (SGA)

4: Robert Schick, McKee

1996: Dan Smith, "Effective Systems Design: Plus ça change, plus ça change" (McKee)
1997: "The Memory Bandwidth Problem" (McKee)
1998: "The Memory Bandwidth Problem" (McKee)
1999: "The Memory Bandwidth Problem" (McKee)
2000: "The Memory Bandwidth Problem" (McKee)
2001: "The Memory Bandwidth Problem" (McKee)
2002: "The Memory Bandwidth Problem" (McKee)
2003: "The Memory Bandwidth Problem" (McKee)
2004: "The Memory Bandwidth Problem" (McKee)
2005: "The Memory Bandwidth Problem" (McKee)
2006: "The Memory Bandwidth Problem" (McKee)
2007: "The Memory Bandwidth Problem" (McKee)
2008: "The Memory Bandwidth Problem" (McKee)
2009: "The Memory Bandwidth Problem" (McKee)
2010: "The Memory Bandwidth Problem" (McKee)
2011: "The Memory Bandwidth Problem" (McKee)
2012: "The Memory Bandwidth Problem" (McKee)
2013: "The Memory Bandwidth Problem" (McKee)
2014: "The Memory Bandwidth Problem" (McKee)
2015: "The Memory Bandwidth Problem" (McKee)
2016: "The Memory Bandwidth Problem" (McKee)
2017: "The Memory Bandwidth Problem" (McKee)
2018: "The Memory Bandwidth Problem" (McKee)
2019: "The Memory Bandwidth Problem" (McKee)
2020: "The Memory Bandwidth Problem" (McKee)
2021: "The Memory Bandwidth Problem" (McKee)
2022: "The Memory Bandwidth Problem" (McKee)
2023: "The Memory Bandwidth Problem" (McKee)
2024: "The Memory Bandwidth Problem" (McKee)
2025: "The Memory Bandwidth Problem" (McKee)

Sally A. McKee

Computer Systems Laboratory

Electrical and Computer Engineering

CORNELL


The Memory Wall

Made simplifying assumptions

- $t_{\text{avg}} = p \times t_{\text{cache}} + (1-p) \times t_{\text{memory}}$
- Every 5th instruction references memory
- CPU speeds increase 50-100% / year
- DRAM speeds increase 7% / year


How long before **ALWAYS** waiting for memory?

The Original Prediction



Since the destruction of the temple,
prophecy has been given to the fools

A Picture's Worth



95% CPU idle
times for HPC
scientific TP
applications

Why?


Lack of reference **locality**

- Registers
- Cache lines (\forall caches)
- TLB entries (btw, TLB == cache)
- VM pages (yup, VM == cache)
- DRAM pages (caching here, too)

Contention for resources

almost dual of locality
optimizations

Non-uniform DRAM Access



Possible Approaches

Use bigger, deeper cache hierarchies

Add more/better latency-tolerating features

- Non-blocking caches
- Out-of-order instruction pipelines
- More speculation
- Multithreading

Migrate intelligence \leftrightarrow DRAMs

Isolates complexity
within one
component

Create smarter memory subsystems 

Make software control how cache is managed

Smarter How?

Know how to build memory efficiently, cost-

effective Prefetch read data

• Buffer write data
scatter/gather REALLY well, makes good

Remap addresses
use of DRAM resources, and makes

• Use cache capacity better
better use of on-chip cache resources
• Use bus bandwidth wisely

Schedule backend (DRAM) accesses better

Won't slow down normal accesses
• Optimize use of memory bus

Won't chance CPU
• Exploit parallelism among DRAM banks

Will perform like SRAM
• Exploit locality in page buffers (hot rows)


Sally A. McKee

Computer Systems Laboratory
Electrical and Computer Engineering

CORNELL

Motivating Example

```
for (i = 0; i < n; i++)  
    sum += A[i][i];
```



Wasted bus bandwidth

Low cache utilization

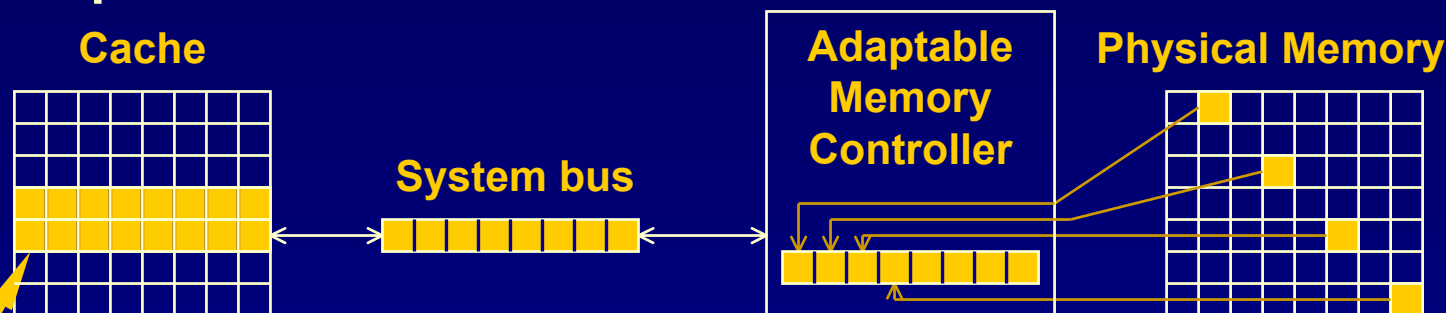
Low cache hit rate

Low TLB hit rate

Gathering within the Impulse MC

Load only data needed by processor

Gather sparse data to dense cache lines




```
diagonal = remap_strided(...);  
for (i = 0; i < n; i++)  
    sum += diagonal[i];
```

20x speedup for
4K×4K array

Impulse Remapping

Exploit unused physical (shadow) addresses
Remap at fine or coarse granularity



Indirection Vector Remapping

```
for (i=0; i<n; i++)  
    ...A[iv[i]]...;  
aA = remap_indirect(..)  
for (i=0; i<n; i++)  
    ...aA[i]...;
```

Memory controller maps $aA[i] \Rightarrow A[iv[i]]$

Indirect accesses replaced by sequential

Accesses to iv moved to MC

Dynamic Indirection Vectors

Don't know entire **iv** ahead?

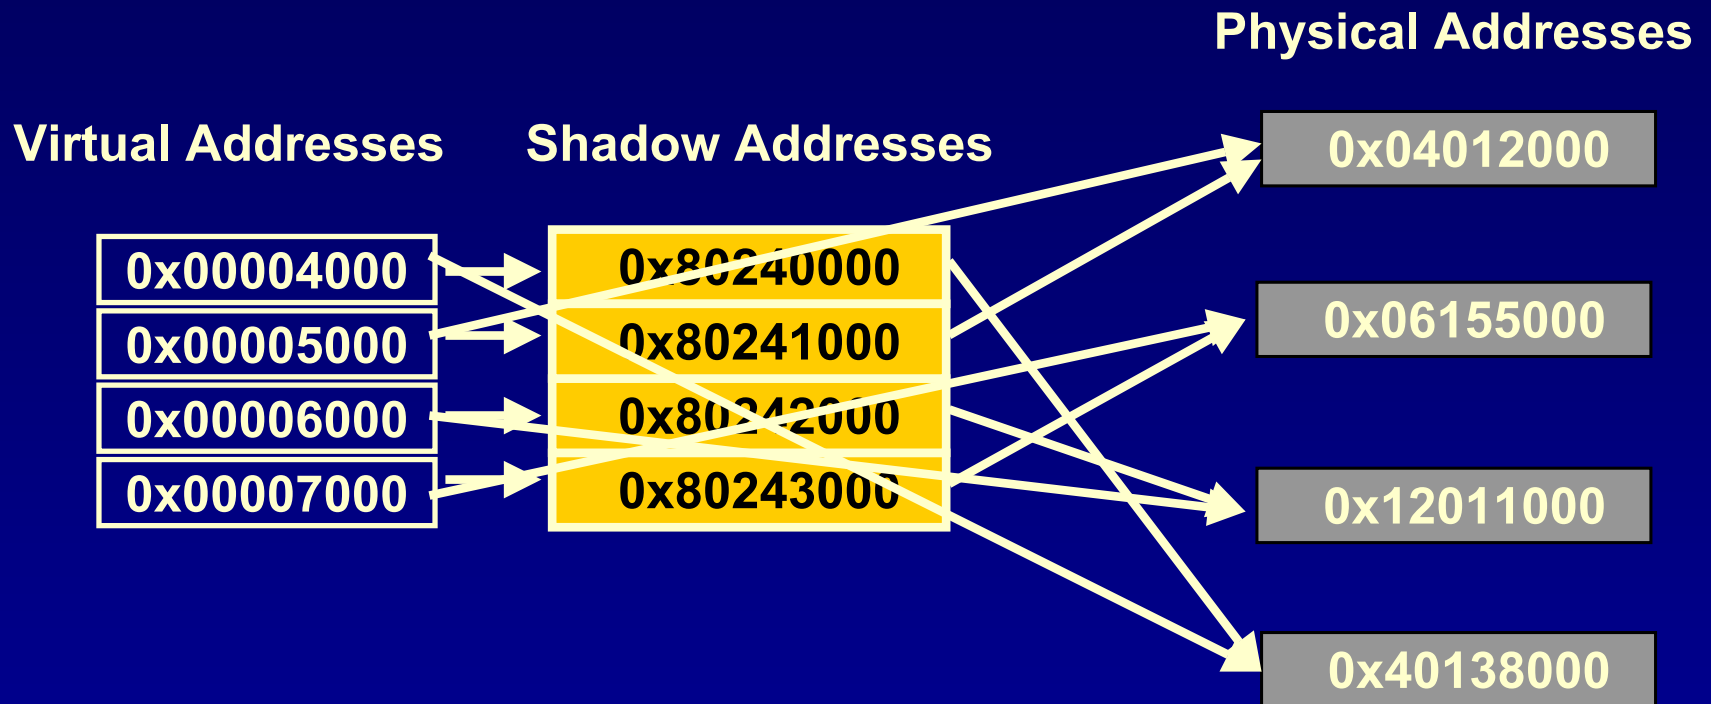
```
for (i=0; i<N; i++)  
    sum += A[random()];
```

Stripmine loop:

```
aA = remap_DIV(A, &iv, 32, ...);  
for (i=0; i<N/32; i++) {  
    for (k=0; k<32; k++)  
        iv[k] = random();  
    flush_to_MC(iv);  
    for (k=0; k<32; k++)  
        sum += aA[k];  
    purge_from_cache(aA);  
}
```

} Analogous
to get/put

No-Copy Superpages



The Impulse “Big Picture”

Improve memory locality via remapping

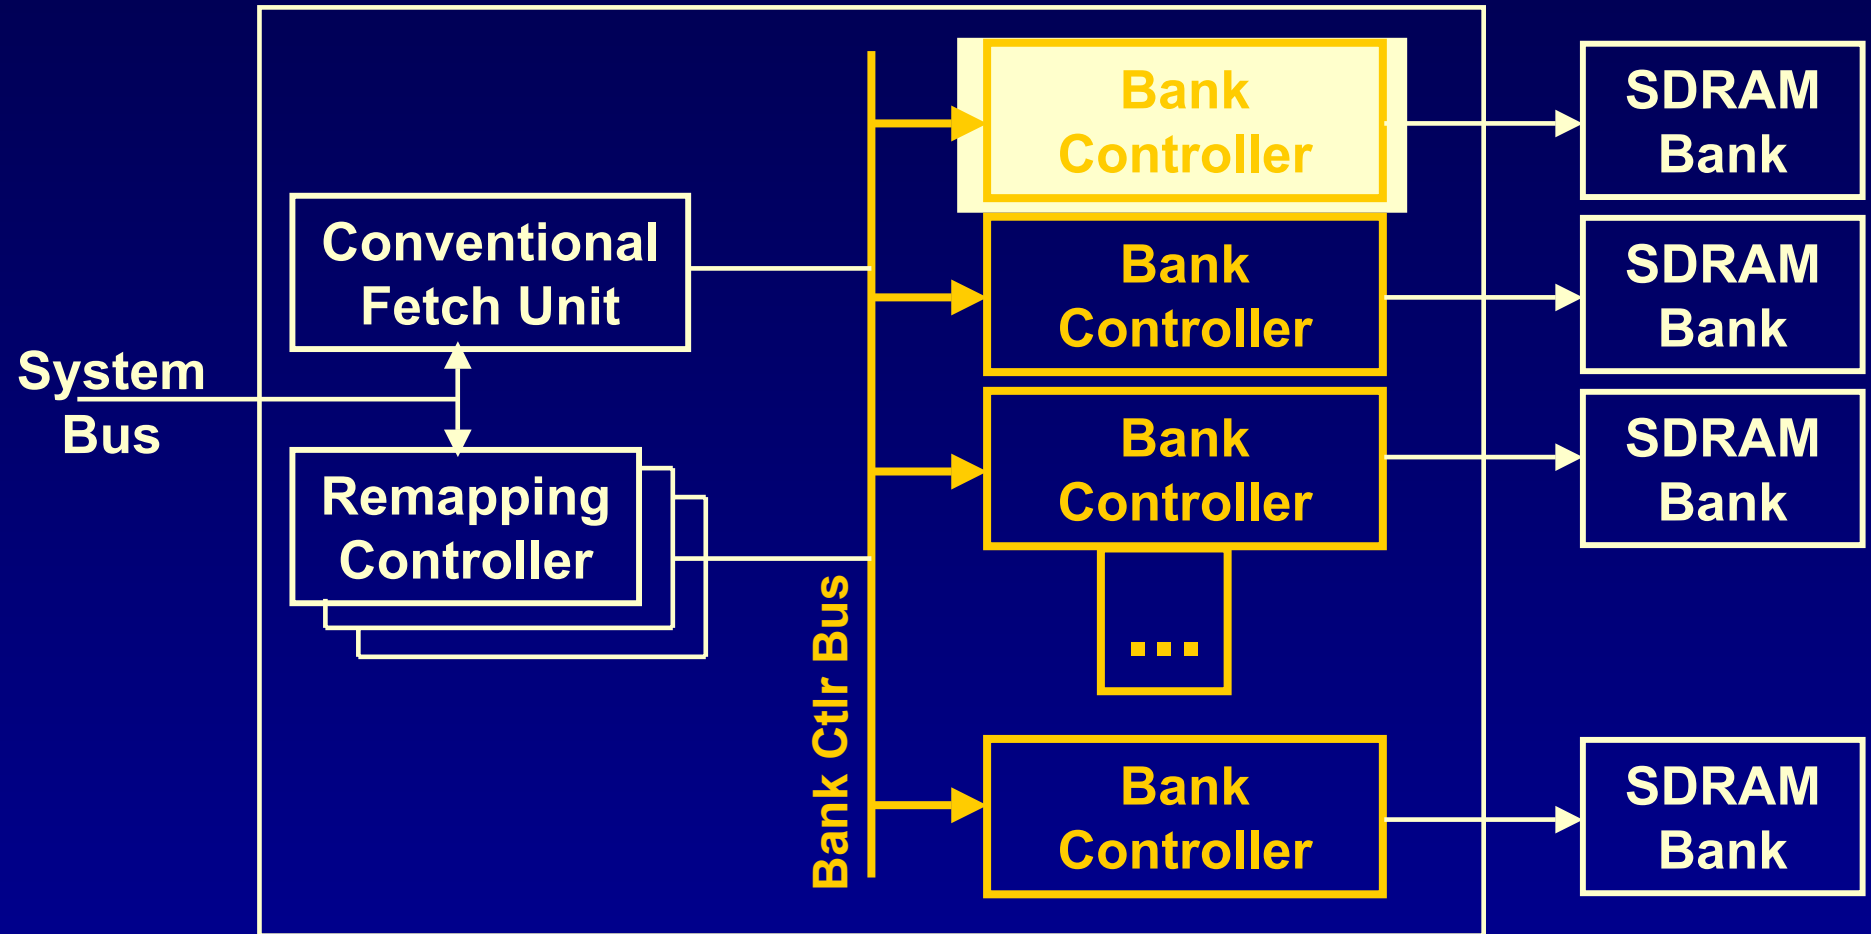
- Improve system bus utilization
- Increase cache efficiency

Increase throughput with parallelism

- Overlap CPU/memory activity
- Exploit parallel SDRAM banks

Exploit SDRAM's NUMA characteristics

Conceptual Organization



Parallel Vector Access Backend

Remapping controllers issue special **vector** ops

- Base-stride: issue (first address, stride, length) tuple
- Vector-indirect: issue four indices per cycle (tentative design)

Bank controllers make independent decisions

- Am I involved in this vector read?
- What elements must I fetch?
- How can I fetch them most efficiently?

When all elements fetched on a read ...

- Control lines indicate completion of vector read
- Coalescing done via wired-OR operations
- Bank controller bus speed matches system bus

PVA Solution Details

$V = \langle V.B, V.S, V.L \rangle$ (base, stride, length)

Fast Basic Functions

- $\text{FirstHit}(V, b)$: first vector element of V that hits b
Table lookup, multiply or shift and add
- $\text{NextHit}(V.S)$: incremental index of next element
Trivial PLA

Bank Controller Algorithm

Compute $i = \text{FirstHit}(V, b)$

If no hit, continue

Repeat until end of the vector:

Schedule access to memory location $V.B + i * V.S$

$i = i + \text{NextHit}(V.S)$

Scheduling Heuristics

- Early row open
- Reordering and interleaving requests


Sally A. McKee

Computer Systems Laboratory
Electrical and Computer Engineering

CORNELL


Cache-line Interleaved, Serial Vector Gathers

$V = \langle 1024, 1, 16 \rangle$ (same as cache-line fill)




Cache-line Interleaved, Serial Strided Vector Gathers

$V = \langle 1024, 32, 16 \rangle$ (vector gather)




Word Interleaved, Serial Vector Gathers

$V = \langle 1024, 1, 16 \rangle$ (cache-line fill)



Word Interleaved, Serial Strided Vector Gathers

$V = \langle 1024, 2, 16 \rangle$



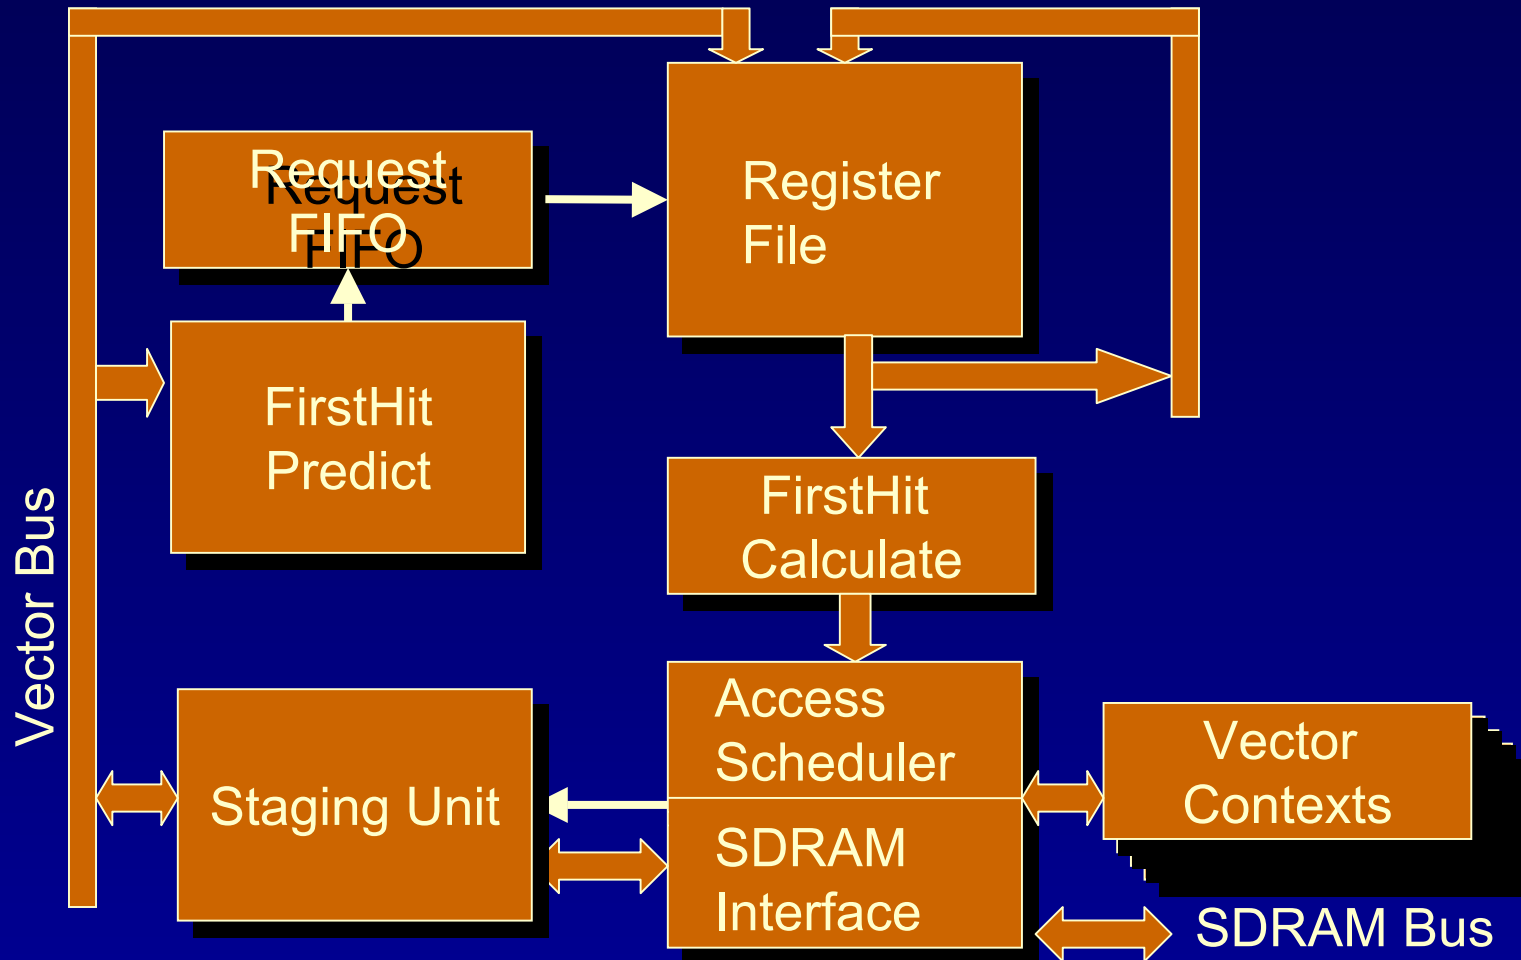
PVA Stride-2 Gather



Bank Controller Components

- Firsthit predictor
- Request FIFO
- Register file
- Firsthit calculator
- Access scheduler
- Vector contexts
- Scheduling policy module
- Staging units (read and write)

Bank Controller Organization




Sally A. McKee

Computer Systems Laboratory
Electrical and Computer Engineering

CORNELL


Stride-1 Vectors (Cache Line Fills)



SDRAM PVA takes about same time as SRAM system

PVA takes about same time as cache-line optimized controller


Stride-2 Vectors



SDRAM PVA takes about same time as SRAM system

PVA takes about 1/2 time of cache-line optimized controller


Stride-4 Vectors



SDRAM PVA takes about same time as SRAM system

PVA takes about 1/3-1/5 time of cache-line optimized controller

Stride-16 Vectors



Can't exploit bank parallelism as well


SDRAM PVA still takes about same time as SRAM system

Sally A. McKee

Computer Systems Laboratory
Electrical and Computer Engineering

CORNELL

Stride-19 Vectors (Diagonal Example)



PVA takes about same time as SRAM memory system

PVA takes about same time as for stride-1 vector

PVA Results Summary

FPGA Synthesis:

- 3600 lines of Verilog
- 10K logic elements and 2K on-chip RAM
- FirstHit() requires 2 cycles (under 20nsec at 100MHz)
- NextHit() requires 1 cycle
- Minimal increase in hardware complexity

Highlights of performance:

- Stride 1: PVA fast as usual cacheline-optimized serial unit (99%-108%)
- Stride 4: PVA 3x faster than pipelined serial gather unit
- Stride 19: PVA up to **33x faster** than cacheline-optimized serial unit
- Specific gains depend on relative skew of the various vectors
- 2-5x faster than similar proposed designs

The Bad News

These are uniprocessor solutions

- Working on SMP adaptations
- Require hardware/software changes
- Complexity still isolated

Have to restructure code

- Compiler can do much of the work
- Can semi-automate the rest?
- **Need better tools**

Tools Wish List

Memory performance monitoring

- Better metrics
- Automatic identification of bottlenecks

Visualization

Interactive performance tuning

- Let compiler do what it can
- Exploit user's knowledge of application
- Exploit **temporal** locality better

So What Do We Do?

We're stuck with DRAM

- Economics
- Lack of viable alternatives

Everything we can

- Change hardware (where possible)
- Restructure code (at least recompile)
- Build better tools

The Impulse Team

- John Carter
- Al Davis
- Wilson Hsieh
- Kathryn McKinley
- Binu Mathew
- Mike Parker
- Lixin Zhang
- Zhen Fang
- Ali Ibrahim
- EE masters students

Questions?

www.cs.utah.edu/impulse

www.csl.cornell.edu/~sam/papers

sam@csl.cornell.edu